Engineering Applications of Artificial Intelligence 27 (2014) 17–27

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Multi-objective evolutionary design of robust controllers on the grid

Alex Shenfield [a,*], Peter J. Fleming [b]

[a] School of Engineering, Manchester Metropolitan University, Manchester, M1 5JD UK
[b] Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield S1 3JD, UK

ARTICLE INFO

Article history:
Received 18 July 2012
Received in revised form
29 March 2013
Accepted 24 September 2013
Available online 14 October 2013

Keywords:
Aerospace stability control
H_∞ controller design
Multiobjective evolutionary algorithms
High performance computing

ABSTRACT

Coupling conventional controller design methods, model based controller synthesis and simulation, and multi-objective evolutionary optimisation methods frequently results in an extremely computationally expensive design process. However, the emerging paradigm of grid computing provides a powerful platform for the solution of such problems by providing transparent access to large-scale distributed high-performance compute resources. As well as substantially speeding up the time taken to find a single controller design satisfying a set of performance requirements this grid-enabled design process allows a designer to effectively explore the solution space of potential candidate solutions. An example of this is in the multi-objective evolutionary design of robust controllers, where each candidate controller design has to be synthesised and the resulting performance of the compensated system evaluated by computer simulation. This paper introduces a grid-enabled framework for the multi-objective optimisation of computationally expensive problems which will then be demonstrated using and example of the multi-objective evolutionary design of a robust lateral stability controller for a real-world aircraft using H_∞ loop shaping.

1. Introduction

Modern aircraft consist of many complex subsystems, all of which require robust and reliable control. These systems are often multi-variable, consisting of multiple inputs and multiple outputs, and frequently the desired responses of a subsystem are in conflict with each other (for example, a controller design that achieves the minimum possible overshoot of the plant often requires accepting a slower rise time than might otherwise have been achieved).

Whilst conventional robust controller design methods such as H_∞ or LQG control can be effectively used to create controllers that are robust both to modelling uncertainties and to cross-coupling between channels in complex multi-variable systems, the resulting controlled system often performs unsatisfactorily. One approach to overcoming this problem is by coupling novel evolutionary multi-objective optimisation techniques with these conventional controller design methods. This provides the engineer with a set of powerful tools for addressing complex multi-variable problems with performance constraints (Fleming and Purshouse, 2002). This type of integrated multi-objective optimisation approach to the design of robust controllers has been successfully used for the design of fixed structure robust H_∞ controllers (Wang and Li, 2011), as well as forming the basis of a novel multi-objective PID controller design procedure (Reynoso-Meza et al., 2012, 2013). However, such methods are frequently computationally expensive, requiring many thousands of controller designs to be evaluated.

Grid computing offers one potential solution to the computationally expensive nature of this evolutionary controller design process. The grid computing paradigm aims to provide “a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities” (Foster and Kesselman, 1999). This paradigm is differentiated from traditional approaches to distributed computing by its emphasis on providing “a seamless, integrated computational and collaborative environment” (Baker et al., 2002) for the solution of complex problems by allowing coordinated resource sharing across dynamic virtual organisations (Foster et al., 2001). By coupling evolutionary multi-objective optimisation techniques with the large scale distributed high performance computing resources offered by the grid computing paradigm, engineers and designers can effectively address many complex, computationally expensive multi-variable problems – including those that require the synthesis of robust controllers as part of the evaluation process. Grid-enabled optimisation of single objective engineering design problems has been successfully integrated into both computer aided engineering workflows (Weng et al., 2012) and multi-disciplinary design workflows (Lee et al., 2009) to provide easy access to powerful real-time analysis and optimization routines. This allows a potential reduction in both design cycle times

* Corresponding author. Tel.: +44 1612476235.
E-mail addresses: a.shenfield@mmu.ac.uk (A. Shenfield), p.fleming@sheffield.ac.uk (P.J. Fleming).

0952-1976/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.engappai.2013.09.015

and development costs, with a commensurate improvement in product quality.

The purpose of this paper is to describe a grid-enabled framework for evolutionary multi-objective design. This framework will then be applied to the design of a robust controller for the flight dynamics of a real-world aircraft – a complex problem with many (often conflicting) objectives to consider. The paper is organised as follows: Section 2 will provide a brief introduction to evolutionary algorithms, their use in control systems engineering and their application to multi-objective optimisation problems; Section 3 will describe the grid computing paradigm in detail and highlight some of the key features that are used in the development of the optimisation framework; Section 4 will discuss the implementation of the grid-enabled framework for multi-objective evolutionary optimisation; Section 5 will demonstrate the use of the grid-enabled optimisation framework in designing robust control systems for the lateral stability of aircraft; and Section 6 will present our conclusions and outline some ideas for further work.

2. Multi-objective evolutionary algorithms

2.1. Background to evolutionary algorithms

Evolutionary Algorithms (EAs) utilise some of the concepts behind natural selection and population genetics to iteratively evolve a population of candidate solutions to a problem (Goldberg, 1989). They both explore the solution space of a problem (by using variation operators such as mutation and recombination) and exploit valuable information present in the previous generation of candidate solutions (by using a selection operator). The trade-off between exploration of undiscovered regions of the solution space and exploitation of promising areas already discovered by the algorithm is extremely important: too much exploration and the algorithm will take too long to converge on a useful solution, too much exploitation and the algorithm may converge prematurely to local optima. Obtaining a correct balance between exploration of the solution space and exploitation of promising solutions is somewhat of a "black art" (Purshouse, 2003), with little guidance available in the literature on setting the parameters that control this balance. Some promising results have been obtained using computational steering frameworks to allow these parameters to be altered during the run-time of the algorithm (Bullock et al., 2002; Shenfield et al., 2007), but this can be time-intensive and requires the engineer to have a good knowledge of both the optimisation problem and the algorithm design. Another potential solution is to use some kind of *self-adaptation* to dynamically change the balance between exploration and exploitation as the algorithm runs (Beyer, 1995; Igel et al., 2007).

One of the main reasons evolutionary algorithms are applicable across many different problem domains (including those where conventional optimisation techniques may struggle) is their use of evaluation function information directly, rather than derivative information or other auxiliary knowledge. For many non-trivial real-world applications this evaluation function information is obtained by computer simulation of the system. For example, in the optimisation of maintenance schedules for gas turbine aero-engines (Shenfield et al., 2010), the cost information for each schedule is obtained by the computer simulation of a candidate solution over a time period of 25 years. However, this use of computer simulation to obtain evaluation function information leads to some additional problems. To ensure that the results gained from the evolutionary algorithm accurately represent the real-world system, the simulation must be complex enough to capture all the relevant dynamics of the true system. Assuming that this level of complexity is obtainable, this can often lead to the simulation becoming very computationally expensive. Since EAs are both iterative and population based, the simulation may have to be run several thousand times which increases the computational requirements (in terms of computer clock cycles) of the optimisation process significantly.

2.2. Multi-objective evolutionary algorithms

Many real-world engineering problems involve the satisfaction of several, often conflicting, objectives. The general form of a multi-objective optimisation problem can be characterised by a vector of objective functions, f, and the corresponding set of decision variables, x, as (note that minimisation can be assumed here with no loss of generality)

$$\min_{f}(x) = (f_1(x), \ldots, f_n(x)) \quad (1)$$

In this case it is unlikely that a single optimal solution will exist. Instead, the solution of this kind of multi-objective problem leads to a set of Pareto optimal points, where any improvement in one objective will lead to a deterioration in one or more of the other objectives.

A set of non-dominated solutions[1] generated by a multi-objective optimisation algorithm is known as an *approximation set* (Zitzler et al., 2003) and the quality of this set can be characterised by three main performance indicators (Purshouse, 2003):

- The **proximity** of the approximation set to the true Pareto front.
- The **diversity** of the distribution of solutions in the approximation set.
- The **pertinency** of the solutions in the approximation set to the decision maker.

These concepts are illustrated graphically in Fig. 1, where it can be seen that the ideal approximation set produced by an optimiser should be both as close as possible to the true Pareto front (i.e. having good proximity) and provide a uniform spread of solutions across the region of interest of the decision maker (i.e. having a diverse set of candidate solutions that are pertinent to the decision maker).

Conventional multi-objective optimisation methods (such as the weighted sum method Hwang and Masud, 1979 and the goal attainment method Gembicki, 1974) often struggle to satisfy these requirements in the optimisation of real-world engineering problems as they can find only a single point from the approximation set rather than a diverse distribution of potential solutions. This means that a decision maker cannot fully understand the shape of the trade-off space (and thus know whether the *a priori* trade-offs they have chosen are appropriate) without running the optimisation routine many times. However, since evolutionary algorithms search a population of candidate solutions in parallel, they are able to find multiple non-dominated solutions from this approximation set. This provides the decision maker with a set of potential solutions to choose from, rather than a single solution that may not meet the required performance criteria.

A further complication in the application of optimisation routines in real-world engineering design problems is that the optimiser is often required to deal with a large number of objectives. This has led to interest amongst the research community in the area of *many-objective optimisation*.[2] The increased scale of many-objective

[1] A solution in non-dominated if there exists no other solution in the set of current candidate solutions that is better in all objectives.

[2] The phrase *many-objective* has been suggested in the Operations Research (OR) community to refer to problems with more than the standard two or three objectives (Farina and Amato, 2004).